# Real-time Composition as Performance Ecosystem

ARNE EIGENFELDT

School for the Contemporary Arts, Simon Fraser University, Vancouver, Canada
E-mail: arne_e@sfu.ca

**This article proposes that real-time composition can be considered a new performance ecosystem. Rather than an extension of electroacoustic instruments that are used within improvisatory environments, real-time composition systems are produced by composers interested in gestural interactions between musical agents, with or without real-time control. They are a subclass of interactive systems, specifically a genre of interactive composition systems that share compositional control between composer and system. Designing the complexity of interactions between agents is a compositional act, and its outcomes are realised during performance – more so than most interactive systems, the new performance ecosystem is compositional in nature.**

## INTRODUCTION

In his 2007 article on performance ecosystems, Waters, in describing Adam Green's 'Suit' for performance, suggests a 'portable performing environment is not unusual as a compositional motivation' (Waters 2007). In a complementary manner, I would suggest that my notion of real-time composition (Eigenfeldt 2008a) offers a portable composition environment, arising from a performance motivation. As I am trained as a composer, I view almost every musical situation as a compositional endeavour; admittedly, this might be seen as a step backwards in terms of musicking (Small 1998). While I find immense intellectual pleasure in controlling and defining musical complexity – as I assume many, if not most, composers do – I also feel that this is not to be undertaken at the expense of performance expressivity. My discomfort with presenting fixed media compositions is perhaps related to my appreciation for performance; the excitement and enjoyment of seeing and hearing music created before my very eyes and ears has always drawn me to live electroacoustic music (EA), specifically interactive systems.

However, I do not consider myself a performer, and certainly not an improvisor, and thus have never naturally embraced the performative aspects of live EA. Waters, in describing his relationship to the laptop as compared to the baroque flute – with both of which he performs – explains his ideological position as resulting from his role as performer (Waters 2007). Conversely, my ideological position arises from my role as a composer, interested in the potential for compositional control over multiple complex gestures during performance.

In the Western art-music tradition, the composer has been removed from the performance ecosystem: the act of composition separated from its presentation, and audiences (and performers) afforded the security of knowing that every note has been considered, and would not – indeed, could not – be altered. The initial efforts of electroacoustic art music coincided with the height of modernism and the veneration of such ideals, and it is not surprising that the fixity of acousmatic music continues this tradition. This demarcation still exists today, as young composers approaching electroacoustic music are forced to make a decision: those interested in controlled and defined complexity will favour the fixity of acousmatic or soundscape composition, while those that are interested in more conceptual or mercurial aspects (including musicking) will be drawn to the fluidity of live EA.

Composers who also view themselves as performers may gravitate towards the potential for improvisation within live EA. A standard definition of improvisation, one given by *Grove's Dictionary of Music and Musicians*, is the creation of a musical work as it is being performed. While I doubt *Grove's* considered the implications of interactive systems in relation to composition and improvisation, Benson (2003) gives us a more useable frame of reference, when he gives eleven shades of improvisation, ranging from subtle levels of interpretation to free extemporaneous creation during performance. Interestingly, he includes nonperformative aspects of variation within fixed composition, giving Copland's use of 'Tis A Gift to Be Simple' as an example.

Real-time composition (RTC) is an attempt to unite these two streams, bringing (or perhaps returning) performative aspects to composition, and the control of complexity to improvisation. As an RTC is, by its very nature, an interactive system, it already blurs the distinctions between performer, instrument and environment. However, it adds to this triumvirate the role of composer (while potentially removing that of performer). As such, it raises the

*Organised Sound* **16**(2): 145–153 © Cambridge University Press, 2011. doi:10.1017/S1355771811000094

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

question: is it possible for a composer to be active within the performance ecosystem *without* being a performer?

## 2. REAL-TIME COMPOSITION

Real-time composition is the application of musical agents to interact in musical ways, during performance. Each agent has the potential to control an independent musical gesture – either pitch-based or timbral – and the complexity of the interactions, along with the quantity of simultaneous gestures, cannot be controlled in any detailed way using existing performative actions. In other words, knowledge of how to interact musically must be built into the agents, and an environment created in which these agent interactions can result in artistically interesting and compositionally satisfying sonic artworks.

Real-time composition is not improvisation, just as improvisation is not RTC (Lewis 2000). Although RTC has evolved from improvisatory interactive systems, the complexity desired by composers in RTC cannot be controlled through existing performative methods used in improvisational systems, nor through constrained random procedures (Eigenfeldt 2007). Furthermore, limits to human dexterity (the physical inability to control more than a limited number of musical gestures at once), interaction and prediction (performers will not know exactly what others intend to do, thereby limiting the potential for goal-driven motion), and memory (the inability to incorporate direct or indirect repetition of earlier musical material, a standard element of musical form) are elements addressed in RTC that differentiate improvisation from composition. Imbuing multi-agents with musical knowledge and intelligence, and facilitating their interaction in real-time, allows for the creation of compositional environments during performance. Real-time composition is, first and foremost, a compositional medium, albeit one that is based within performance.

Using Drummond's taxonomy of interactive systems (2009), RTC is situated within interactive composing, comprising a system in which the musical output is created through shared control between the system's programming and potential performance input. An RTC system, like Chadabe's interactive composition system (Chadabe 1984), can be designed to create a single composition that may change substantially from performance to performance. The result is a further blurring of the traditional performance ecosystem between instrument, composer and performer: 'The instrument is the music. The composer is the performer' (Chadabe 1997). The difference between interactive composing systems and RTC is the latter's emphasis on the intelligence and autonomy of the system.

## 3. COMPOSITION, IMPROVISATION AND CONTROL

Waters suggests that the separation between composition and improvisation has been a recent phenomenon, and that current developments (mainly due to changes and advances in technology) will convolve the two within music as practice (Waters 2007). From my perspective within academia and art-music circles – which is in no way representative of the wider music-making population – I am not so certain. While many artists are traversing both ends of the fixed versus fluid spectrum, and exploring the continuum between, the fixed-time object is an ideal that many composers will continue to find attractive.

Lewis sees the critical distinction between composition and improvisation, a separation he ascribes to the Euro-American art-music culture and fostered in academic institutions, in framing improvisation as involving 'effortless spontaneity' versus the 'careful deliberation of composition', somewhat facetiously summarising this opposition as 'composer as ant, the improvisor as grasshopper' (Lewis 2000). Lewis suggests that a distinction between composition and improvisation rests on the 'axis of interaction', in that improvisation is open to contingency. The 'dialogic imagination' of improvisation occurs in real-time; however, contingency and feedback naturally exist during the compositional stage of fixed media works, even if it is assumed to be lost in performance.[1] He does acknowledge that his own interactive system, *Voyager*, is a 'composing machine that allows outside intervention' (Lewis 2000: 38).

Interactive systems have challenged and extended the notion of instrument, specifically in regards to performance control. Traditional acoustic instruments imply a significant degree of control, while interactive systems introduce a shared control between performer and system that creates an interaction that 'may well be considered more in terms of a composition or structured improvisation rather than an instrument' (Drummond 2009: 124). How each system is framed – improvisation, instrument or composition – reflects the particular performance ecosystem in which the designer is interested. For example, Lewis' *Voyager* (and many other performer plus system models) allows the performer to interact with the system through acoustic input, as if it were another improviser; Waisvisz's *The Hands* allows the performer to guide a quasi-predictable instrument using gestural controllers; Chadabe's *Solo* is an example of a composition created interactively. Real-time composition situates itself in this latter family.

[1]Fluidity and fixity thus occur in different temporal stages in the life of a musical work: in RTC, the fluidity is evident during performance and the fixity exists within the system itself, whereas, in non-RTC, fluidity occurs during the compositional (or pre-compositional) stage, and the fixity in the artefact.

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

### 3.1. Complex systems

While distinctions between performer, instrument and environment are blurred within interactive systems, the notion of feedback remains integral. The computer – whether it is considered an instrument or performer – is expected to make decisions regarding the input it receives, decisions that result in a variable sonic output. Such decisions are typical of improvisational participants; as Casal suggests, 'improvisation happens in an environment full of snap judgments, where previous experience, cultural heritage, and current information acquired through listening all help enable the improvisor to make decisions quickly' (Casal and Morelli 2007: 3). If the performance ecosystem expands beyond a single performer and instrument, to modelling the complex interactions of a collective – Lewis' *Voyager*, for example – the nature of the model will be crucial to its success. Earlier interactive systems, such as those by Lewis and Chadabe, relied upon constrained random procedures (Winkler 1995), a limitation that I discuss more fully elsewhere (Eigenfeldt 2007).

Newer systems have explored complex systems within performance. Casal has chosen to use elements of Artificial Life (A-Life) and Music Information Retrieval (MIR) within his system *Frank*. Impett (2001) has implemented a swarm system, which he influences in performance with his extended trumpet. Swarm systems also have been musically realised by Blackwell and Young (2004). Young (2008) and Bown and Lexer (2006) have implemented neural networks that learn during performance; A-Life strategies have been employed by composers, including Biles (1994), Beyls (2009) and Weinberg, Godfrey, Rae and Rhoads (2008). Designing complex systems and influencing those systems in performance – thereby designing the interaction as well – moves beyond the improvisational to the compositional; however, the nature of the complexity is a defining aspect. Within the above-mentioned examples, the complexity of these systems is *organised* (Weaver 1948) – as it is achieved through the interaction of its parts (and has the potential for emergent properties) – whereas the complexity of earlier constrained random systems' is *disorganised*, as it is generated (and can be understood) by statistical models (i.e. random probabilities). As Impett and others have demonstrated, controlling an organised complex system in performance is feasible.

While the individual elements within such systems may behave in complex and unpredictable ways, their macro-level interactions are governed by a defining principle. In designing such interactions, we are thus embracing Di Scipio's reformulation of interactivity as 'a by-product of carefully planned-out interdependencies among system components' (Di Scipio 2003: 27). Such complexities of interaction emerge from free improvising human ensembles – however, the organisation is collective; if organisational methods are defined by an individual, it is an act of composition. Therefore, if singular (as opposed to collective) and significant (as opposed to stochastic) control over complexity can exist within live performance, this new performance ecosystem is compositional, and can be considered real-time composition.

## 4. TOWARDS A REAL-TIME COMPOSITION SYSTEM

A great deal of research has gone into modelling human improvisation (see Collins 2006 for a detailed discussion). Suffice to say, modelling creative and intelligent human interaction within a performance situation has proven to be extremely difficult: any decision a performer makes, in reaction to a wealth of information arriving through interactions with other performers, will be influenced by not only the given performance conditions – including, but not limited to the current composition's evolution – but also the performer's musical training, musical history, and specific historical interactions with the other performers on stage. Rowe (1992) (quoted in Winkler 1995) posits a view that a software improvisor should be modelled upon human activity: 'interactive software simulates intelligent behaviour by modeling human hearing, understanding, and response'. While others have pointed out limitations to this particular model (Paine 2002; Jordà 2005), many positive developments in real-time musical learning have recently been undertaken (Pachet 2003; Bown and Lexer 2006; Young 2008; Beyls 2009; Maxwell, Pasquier and Eigenfeldt 2009). Despite these exciting advances, a cognitively modelled software improvisor that is successful at both the scientific and artistic levels is, unfortunately, some years away.

### 4.1. Multi-agents

Rather than model human musicians, some systems have modelled the interactions *between* musicians: multi-agent systems have provided some exciting developments in this regard. Intelligent agents are elements of code (programs or patches) that operate without direct user interaction (they are autonomous), interact with one another (they are social), interact with their environment (they are reactive), and make decisions as to when they should operate, and what they should do (they are proactive) (Wooldridge 2009). Since these are also attributes required of musicians in improvisational settings, the use of agents to emulate human–performer interaction has proven to be a fertile field of research. Whalley (2009) gives an overview of the current state of software agents in music and sound art.

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

### 4.2. Artificial Life

Musical agents can exist within a variety of contexts: one that relates to performance ecosystems and offers potentially exciting new models is Artificial Life. A-Life is a discipline that studies natural living systems by simulating some of their biological aspects on computers (Langton 1997), and has subsequently fostered the development of a pool of research tools for studying complexity, including modelling tools for interacting musical agents (Miranda and Todd 2003). See Whitelaw (2004) for an excellent discussion on art and artificial life.

### 4.3. Sonic ecosystems

Bown (2009a) discusses failures of the traditional A-Life paradigm of the interactive genetic algorithm for artistic creation, and points to new approaches based within social learning, cultural dynamics and niche construction that offer potential solutions (McCormack and Bown 2009). Rather than exerting direct control over evolving systems through aesthetic selection, Bown suggests that an ecosystem approach 'might generate sonic works that continue to develop and transform indefinitely, but with consistent structural and aesthetic properties'. Instead of designing methods of selection, an ecosystem approach allows artists to 'design complex generative dynamics in an exploratory manner' – in effect, designing the environment and the way in which sound-making individuals interact, so as to produce interesting artistic results. The significant contribution of this approach is the interactive nature of the design – which allows artists to 'tinker' with the system until it produces desirable outcomes – as well as the fact that the entire system is conceptualised within its medium: a sonic ecosystem operates in the medium of sound rather than being a sonification of a process. *RiverWave* (McCormack and Bown 2009) is an example installation that demonstrates their concept of a sonic ecosystem.

Bown describes a system (2009b) which 'animates the sounds of a man-made environment, establishing a digital wilderness into which these sounds are "released" and allowed to evolve interdependent relationships'. The analysed acoustic environment of the installation space, which includes the projected sounds of the agents, is treated as the virtual environment in which the agents exist. A continuous sonogram is made of the acoustic space, and agents attempt to inhibit low-energy regions, thereby creating an interesting dilemma: as the agents make sound at their occupied frequency region, they deplete the available resources at that region. It is the interaction of agents competing for fitness and available resources that creates the potential for emergence, and thus provides the musical evolution.

The sonic ecosystem (SE) is an exciting RTC model, as it allows for the design of a complex system that serves a musical purpose. Although interaction is not inherent to SEs, it is not excluded.

## 5. REAL-TIME COMPOSITION AS A PERFORMANCE ECOSYSTEM

The author has described elsewhere (Eigenfeldt 2008a) that imbuing multi-agents with musical knowledge and intelligence, and facilitating their interaction in real-time, allows for the creation of compositional environments during performance, and therefore a compositional performance ecosystem. Pre-performance compositional activities, such as defining the environment and the capabilities of the musical agents, can be balanced with non-deterministic qualities of generative music. One example of RTC that combines interactive performance with intelligent agents is the author's *In Equilibrio*, described in detail elsewhere (Eigenfeldt 2009a). In this work, agents communicate and adjust their actions based upon their understanding of the musical environment; performance control is restricted to influencing agent parameters (for example, of pitch centroid), as well as altering the environment itself.

Real-time composition environments are closely related to McCormack and Bown's notion of a sonic ecosystem. They are designed with musical interaction as their first criterion, with sound, in and of itself, as their media. In every case, the systems are performative, in that their output is produced over time, rather than collected for fixed media composition. While some RTC systems have incorporated 'pure' A-Life components of breeding, mutation and selection (Eigenfeldt 2009b), such practices are not necessary within SEs, nor RTCs.

### 5.1. Gestural interactions and plans

One important aspect of composition is the detailed specification of the interaction *between* musical gestures. When serendipitous relationships occur between performers during improvisation, they can appear magical; within fixed compositions, such correlations are predetermined (although potentially still appearing as magical to the listener). These interactions can be incorporated into the design of RTC systems through the use of *plans*, an interesting capability of agents using the Belief-Desire-Intention model (Rao and Georgeff 1995). This model expands on the social attribute of agents, in which they form beliefs about their environments, individual desires and goals, and plans to be executed. Each of these attributes are continually adjusted in performance (thereby offering further parallels between the virtual and the physical performer); however, the ability to share (and adjust) these plans moves beyond human capabilities.

Preliminary research has been undertaken to codify the gestural interactions within acousmatic music

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

(Basanta and Eigenfeldt 2010) and mixed-media works (O'Callaghan and Eigenfeldt 2010) in order to incorporate such potential relationships within future RTC systems.

### 5.2. A preliminary system: kinetic engine

The creation of a true RTC system has been an ongoing investigation by the author, and grew out of work in improvisational interactive performance systems. The first research project created *Kinetic Engine* (KE): its development, both technically and artistically, is discussed elsewhere (Eigenfeldt 2009c). However, some of the salient points can be reframed here within a performance ecosystem.

The first version (KE1) was an attempt to manage high-level decisions within software, generating rhythmic compositions in real-time. While successful in achieving these goals, it also removed the composer from the performance ecosystem, something with which I felt (at the time) somewhat uncomfortable. It was presented as an installation, producing continual short (three- to fifteen-minute) compositions; a great deal of effort was placed in logging the generated parameter settings (i.e. tempo, time signature, timbre, section durations, degrees of variation) of each composition, and fuzzy logic was used to avoid successive iterations from being too similar. Lastly, in an effort to maintain some interactivity, certain performance parameters could be influenced by the audience via sensors. As the first audience comprised computer music creators and developers,[2] I had assumed audience members would attempt to delicately influence the system, listening to the subtle shifts in relationships between agents; I was (perhaps naively) surprised that those audience members bold enough to interact with the system immediately tested the limits of interaction and control, varying parameter settings to their extreme, essentially making the system choke.

Having removed myself from the performance ecosystem, my role as passive observer allowed me to take copious notes on the successes and failures of each short composition, and audience reactions to them. Prior to the next day's presentation, I adjusted many of the parameter ranges so as to maximise success and minimise failures (which can be understood as uninteresting music). This was a curious delayed (or perhaps anticipatory) feedback, which, in the end, resulted in an artistically unsatisfactory performance ecosystem.

The next system, KE2, returned to an interactive model that placed the composer within the performance, using immediate feedback to guide the system at a high level by influencing agent interactions (see Eigenfeldt 2008b for a detailed description). However, KE2 was not a true RTC, but rather an intermediary step between interactive multi-voiced instrument and intelligent composition; the agents had no plans or desires, other than to interact in rhythmically interesting ways.

Both versions of the system raised a design issue concerning the balance between elements of surprise versus guaranteed success. As it was operating under a ruleset, the more defined this set became, the more likely it was to generate music I wanted to hear. At the same time, limiting this rule-set also limited the potential for genuine surprise.[3] During performance, KE2 was a consistent performer that could be relied upon to generate complex and musically interesting rhythms; however, it was not very versatile, and, while its details continued to surprise me, it never went anywhere unexpected.

My solution was to create a new system that allowed not only for performance interaction, but also an adaptive and dynamic ruleset through corpus-based recombinance. *Kinetic Engine* 3 (KE3; Eigenfeldt 2009b) used a genetic algorithm to continually evolve derived rhythmic phrases; another (Eigenfeldt and Pasquier 2010) used a steerable Markov model to generate harmonic progressions. Like all interactive systems, the design of the interaction – between agents, in this case – was a compositional decision; the selection of the corpus to be used was a further compositional act. Performance actions were limited to very high-level choices, such as when to change sections, and influencing agent interactions in global ways. Unfortunately, the amount of calculations required to continually generate an evolving population of individuals in real-time proved to be too demanding for performance, and KE3 was limited to a non-performance role, producing fixed works.

## 6. COMING TOGETHER: A SERIES OF RTC SYSTEMS

More recently, I have avoided the question of balancing novelty and quality in a general-purpose system, and focused upon creating RTC systems that can only produce a single composition whose contents vary with each performance: quality is assumed to exist within the compositional constraints imposed upon the environment, and novelty within the exploration of these environments. In these systems, agents negotiate musical content within a defined musical environment, with or without direct performer interaction. In each case, agents are given random musical material, and through the convergence of predefined musical

[2]KE1 was premiered at the 2006 International Computer Music Conference in New Orleans.

[3]The qualities of novelty and quality are not antithetical, by any means; in fact, they can be viewed as two key properties of computational creative systems (Ritchie 2007).

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

parameters, self-organisation is demonstrated. Each system is situated within different environments: *Beauty & Truth* converges pitch- and timbre-based agents in melodic and harmonic phrases; *Freesound* converges soundscape recordings and real-time signal processing; *Shoals* converges granular synthesis parameters within a spectral analysis database of real-time audio. The encompassing title of *Coming Together* is a reference to the movement from randomness to order that occurs within each system.

### 6.1. *Coming Together: Beauty & Truth*

The first *Coming Together* system is conceived as an autonomous negotiation between melodic agents, which, given no predefined material, converge towards a musically meaningful cumulative output, whose content is generated and determined by the actions and interactions of the agents themselves. The result is a continuous series of compositions (usually between 8 and 15 minutes each) which either end in success – marked by the instantiation of additional agents – or failure, in which the agents collectively acknowledge failure and start over.

An individual agent's desire is to generate a repeating melodic phrase that has a meaningful relationship to all other agent phrases. By musically meaningful, I mean harmonic, melodic and rhythmic associations that could have been produced by a human composer: Ritchie's *value* property (Ritchie 2007). Agents move towards this collective goal by listening to one another's musical output – which can be MIDI data or audio – as well as messaging one another, altering their own musical output based upon formed beliefs as to other agent's states. The end result is a negotiated solution that demonstrates a unique melodic, harmonic and rhythmic environment that has been arrived at collectively.

Constraints are placed upon agent choices to engender reasonable outcomes. Although complete unpredictability may be desirable from a standpoint of A-Life research, the choice was made to use a sonic ecosystem approach in which the environment was designed to produce artistically satisfying results. While the design bears a resemblance to the first KE system, there is a great deal more intelligence in the interaction between the musical agents, resulting in a higher degree of novelty within the musical details; however, the composition's overall structure, a consequence of the converging agent desires, produces the clear notion of a single composition continually re-interpreted.

The performance ecosystem in *Beauty & Truth* is certainly challenged, as several components seem to be missing. Feedback, integral during the development and fine-tuning stage, is neglected during performance. The composition is set in motion, and its outcome cannot be influenced until the system succeeds or fails. While each composition will be different, the system does not take into account any environmental aspects, and will happily perform for an empty room, thereby ignoring Impett's model for interactive music as 'music instantiated in real-time on the basis of local performance and environmental information' (Impett 2001: 108). By this account, one can question whether performance is actually involved, and whether this work is instead simply the sonification of a complex process.

An important design decision in *Beauty & Truth* was the deliberate *avoidance* of performance. Western notions of concert performance imply structures with a beginning, middle, and end, and a predictable duration. Installations, on the other hand, can be viewed as ongoing: generative installations – in which *Beauty & Truth* situates itself – allow for the continued revisitation by an audience, stepping in and out of the ecosystems, so as to listen to its diversification over time. I have found the installation model to often be more suitable to RTC, as it can then operate outside the bounded time constraints of the concert presentation; interestingly, other creators of A-Life systems have similarly chosen to present their work in such a fashion (McCormack 2001; McCormack and Bown 2009; Bown 2009b).

### 6.2. *Coming Together: Freesound*

Rather than manipulate events, as *Beauty and Truth* does, *Freesound* uses soundscape recordings as its source material. Agents communicate their selections from the pre-analysed library of recordings taken from freesound.org by sending metatag data; agents also have the ability to analyse the spectrum of other agents' audio. During performance, agents listen to the sonic environment, and select material from their unique database based upon their perception of the current conditions (which is formed from both spectral analysis and interpretation of the collective metatag data). At different points of the composition, selection methods vary.

An important contribution to RTC design in *Freesound* is the use of a defining structure that is generated at the beginning of the performance. *Freesound* is framed as a performance (rather than installation), and seeks to balance agent negotiation with predefined compositional structure without the necessity of user interaction to guide it. Certain aspects of the environment's evolution in time are predefined:

- Prior to performance, the user indicates an overall duration from which four sections are generated algorithmically. Agents are aware of their current position within the structure, and behave differently depending upon the point in time.

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

- The global evolution of certain parameters use preset tendency masks.
- The overall increase in resonant filtering is a defining audio feature of the composition.
- The initial selection of recordings is from freesound.org.

*Freesound* demonstrates characteristics of a sonic ecosystem in that the environment is carefully calculated, yet the interaction between its components remains complex and unpredictable. Musically, *Freesound* generates surprising and varied compositions. Ironically, the composer's role is limited to adjusting amplitude levels of the four agent's gestures, similar to the performative diffusion of fixed acousmatic works; however, unlike such works, the ecosystem is dynamic and mercurial, the *content* being negotiated by the agents. And while the performative actions cannot influence the direction of the work – the feedback is one-way – they certainly contribute to its success (or failure, if one neglects to listen and react appropriately).

### 6.3. *Coming Together: Shoals*

This RTC is composed of a population of autonomous agents that exist within a sonic ecosystem derived from real-time analysis of live audio. Agents search for food consisting of CataRT unit analyses (Schwarz 2007), within a two-dimensional projection in a descriptor space of spectral flatness and spectral centroid. When an agent finds food, it broadcasts this information to the community; all agents within that agent's neighbourhood – a globally defined radius – can move towards the location in which the food was discovered. Agents become active when they find food, which is then consumed through granulation.

Individual agents are initialised with random synthesis parameters, but communicate these parameters to agents within their social network. Networks are formed initially through proximity (i.e. an agent's neighbourhood), but later through more advanced selection methods, which include food supply and total numbers of agents within the network. Agents within a social network converge their parameters, thereby creating unified grain streams. Separate gestures thus emerge through the self-organisation of the population.

The musical objective of *Shoals* is the continual generation of distinct gestures that are arrived at through self-organisation, which itself is generated through convergence of agent parameters using a social network. Agent interaction is dependent upon the performance environment and the existence of new audio input. *Shoals* has the ability to be configured over a network, with audio and agent parameter data gathered from geographically displaced locations.

Like the previous described RTCs, *Shoals* is a sonic ecosystem whose components have been carefully adjusted so as to produce musically interesting outcomes, while at the same time providing for novelty and surprise, even given the same environmental conditions. However, the importance of the environment to the sonic outcome has been greatly increased, in that live performance can directly influence the agent actions. Similarly, feedback is two-way, with performers able to react to the evolving timbres. The composer's role in the performance ecosystem is, like *Beauty & Truth*, non-interactive, but clearly evident within the design of the interactions themselves.

A question can be raised as to whether any composed improvisation could be considered real-time composition, if a main criterion is defining the complexity of interactions. I believe the answer to be yes, to the extent that the interactions are clearly demarcated. However, within an RTC, the limitations of agent behaviour can be clearly defined, as well as altered, during performance in ways that are not possible with human performers. The desire within RTC is to balance a fixity in compositional choices (inherent within the design of the system itself) with a fluidity in performance possibilities, and it does so through a new performance ecosystem in which composition is the main constituent.

## 7. CONCLUSION

Interactive systems have blurred the traditional relationship between performer, instrument, composition and environment, and, with their basis within performance, have produced variations on the traditional performance ecosystems. Real-time composition systems, developed to explore compositional notions in performance that are beyond improvisational strategies, are a category of interactive systems that focus upon compositional, rather than performative, generation; as such, they constitute a further innovation within performance ecosystems. Real-time composition systems can also be framed as performative sonic ecosystems, in that the careful design of agent interactions and real-time influence are compositional actions that result in varied output with each performance. These systems are conceived to explore music, rather than A-Life, and operate in the medium of sound rather than being a sonification of a process, thus achieving McCormack and Bown's definition of a sonic ecosystem.

The author's exploration of delimited musical environments, as suggested by the ecosystem model, was predicated by the difficulties in generating a cognitively modelled software improvisor that could not only function as a single human within a live performance, but also control and generate complexities beyond those of a human. While the performance

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

ecosystem model for RTC provides a vast array of future possibilities, the eventual goal is true metacreation:[4] endowing machines with creative behaviour. The author, together with other artists and scientists in the fields of artificial intelligence, multiagent systems and cognitive science, is actively engaged in this research project, the results of which should create more complete RTC systems, involving new performance ecosystems.

## Acknowledgements

This research was funded, in part, by a grant from the Canada Council for the Arts, and the Natural Sciences and Engineering Research Council of Canada. The author is indebted to James Maxwell for valuable comments and suggestions.

## REFERENCES

Basanta, A. and Eigenfeldt, A. 2010. Typological Analysis of Gesture Interaction in Acousmatic Music. *Proceedings of the Electronic Music Studies 2010*. http://www.ems-network.org/spip.php?article300.

Benson, B. 2003. *The Improvisation of Musical Dialogue*. Cambridge: Cambridge University Press.

Beyls, P. 2009. Interactive Composing as the Expression of Autonomous Machine Motivations. *Proceedings of the International Computer Music Conference*. Montreal/San Francisco: ICMA, 267–74.

Biles, J. 1994. GenJam: A Genetic Algorithm for Generating Jazz Solos. *Proceedings of the International Computer Music Conference*. San Francisco: ICMA, 131–37.

Blackwell, T. and Young, M. 2004. Swarm Granulator. Applications of Evolutionary Computing. *Proceedings of The 2004 European Workshop on Evolutionary Music and Art*. Coimbra, Portugal: Springer-Verlag, 399–408.

Bown, O. 2009a. A Framework for Eco-System-Based Generative Music. *Proceedings of Sound and Music Computing 2009*. Porto: SMC, 195–200.

Bown, O. 2009b. Eco-System Models for Real-time Generative Music: A Methodology and Framework. *Proceedings of the ICMC 2009*. Montreal/San Francisco: ICMA, 537–40.

Bown, O. and Lexer, S. 2006. Continuous-time Recurrent Neural Networks for Generative and Interactive Musical Performance. *Applications of Evolutionary Computing: Proceedings EvoWorkshops 2006*. LNCS 3907, Berlin: Springer, 652–64.

Casal, D.P. and Morelli, D. 2007. Remembering the Future: Applications of Genetic Co-Evolution in Music Improvisation. *Proceedings of the European Conference on Artificial Life*, Lisbon.

Chadabe, J. 1984. Interactive Composing. *Computer Music Journal* **8**(1): 22–7.

Chadabe, J. 1997. *Electric Sound: The Past and Promise of Electronic Music*. Upper Saddle River, NJ: Prentice Hall.

Collins, N. 2006. *Towards Autonomous Agents for Live Computer Music: Realtime Machine Listening and Interactive Music Systems*. PhD thesis, University of Cambridge.

Di Scipio, A. 2003. 'Sound is the Interface': From Interactive to Ecosystemic Signal Processing. *Organised Sound* **8**(3): 269–77.

Drummond, J. 2009. Understanding Interactive Systems. *Organised Sound* **14**(2): 124–33.

Eigenfeldt, A. 2007. Real-time Composition or Computer Improvisation? A Composer's Search for Intelligent Tools in Interactive Computer Music. *Proceedings of the Electronic Music Studies 2007*. http://www.ems-network.org/IMG/pdf_EigenfeldtEMS07.pdf.

Eigenfeldt, A. 2008a. Intelligent Real-time Composition. *eContact 10.4 Live-electronics-Improvisation-Interactivity in Electroacoustics*, http://cec.concordia.ca/econtact/10_4/.

Eigenfeldt, A. 2008b. Emergent Rhythms through Multi-Agency in Max/MSP. *Computer Music Modeling and Retrieval: Sense of Sounds, Lecture Notes in Computer Science* **4969**: 368–79.

Eigenfeldt, A. 2009a. Multiagency and Realtime Composition: In Equilibrio, *eContact 11.4 Toronto Electroacoustic Symposium*, http://cec.concordia.ca/econtact/11_4.

Eigenfeldt, A. 2009b. The Evolution of Evolutionary Software: Intelligent Rhythm Generation in Kinetic Engine. *Applications of Evolutionary Computing, EvoWorkshops 2009 Proceedings*, LNCS 5484. Springer: Berlin, 498–507.

Eigenfeldt, A. 2009c. A Realtime Generative Music System using Autonomous Melody, Harmony, and Rhythm Agents. *Proceedings of the 12 Generative Art Conference GA2009,* Milan. Available online at http://www.generativeart.com/on/cic/GA2009Papers/p7.pdf.

Eigenfeldt, A. and Pasquier, P. 2010. Realtime Generation of Harmonic Progressions Using Controlled Markov Selection. *Proceedings of the 1st International Conference on Computational Creativity (ICCCX)*. Lisbon: ACM Press, 16–25.

Impett, J. 2001. Interaction, Simulation and Invention: A Model for Interactive Music. In E. Bilotta, E. Miranda, P. Pantano and P. Todd (eds), *Artificial Life Models for Music Applications*. Consenza: Editoriale Bios, 108–19.

Jordà, S. 2005. *Digital Lutherie: Crafting Musical Computers for New Musics' Performance and Improvisation*. PhD dissertaion, Universitat Pompeu Fabra, Barcelona.

Langton, C.G. (ed.) 1997. *Artificial Life: An Overview*. Cambridge, MA: MIT Press.

Lewis, G. 2000. Too Many Notes: Computers, Complexity and Culture in *Voyager*. *Leonardo Music Journal* **10**: 33–39.

Maxwell, J., Pasquier, P. and Eigenfeldt, A. 2009. Hierarchical Sequential Memory for Music: A Cognitive Model. *Proceedings of the International Society of Music Information Retrieval Conference*, Kobe. http://ismir2009.ismir.net/proceedings/PS3-7.pdf.

McCormack, J. 2001. Eden: An Evolutionary Sonic Ecosystem. *Advances In Artificial Life*, LNCS. Berlin: Springer, 133–42.

McCormack, J. and Bown, O. 2009. Life's What You Make: Niche Construction and Evolutionary Art.

[4]www.metacreation.net.

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press

*Applications of Evolutionary Computing: EvoWorkshops 2009*. *Lecture Notes in Computer Science* **5484**: 528–37.

Miranda, E.R. and Todd, P. 2003. A-Life and Musical Composition: A Brief Survey. *Proceedings of the IX Brazilian Symposium on Computer Music*. Campinas, Brazil.

O'Callaghan, J. and Eigenfeldt, A. 2010. Gesture Transformation through Electronics in the Music of Kaija Saariaho. *Proceedings of the Electronic Music Studies 2010*. http://www.ems-network.org/spip.php?article308.

Pachet, F. 2003. The Continuator: Musical Interaction With Style. *Journal of New Music Research* **32**(3): 333–41.

Paine, G. 2002. Interactivity, Where to from Here? *Organised Sound* **7**(3): 295–304.

Rao, A. and Georgeff, M. 1995. BDI Agents: From Theory to Practice. *Proceedings of the First International Conference on Multi-agent Systems*. San Francisco: AAAI Press, 312–9.

Ritchie, G. 2007. Some Empirical Criteria for Attributing Creativity to a Computer Program. *Minds and Machines* **17**: 67–99.

Rowe, R. 1992. *Interactive Music Systems*. Cambridge, MA: MIT Press.

Schwarz, D. 2007. Corpus-Based Concatenative Synthesis. *IEEE Signal Processing Magazine* **24**(2): 92–104.

Small, C. 1998. *Musicking: The Meanings of Performing and Listening*. Hanover, NH: Wesleyan University Press.

Waters, Simon. 2007. Performance Ecosystems: Ecological Approaches to Musical Interaction'. *Proceedings of the Electroacoustic Music Studies Network 2007*. http://www.ems-network.org/spip.php?article278.

Weaver, W. 1948. Science and Complexity. *American Scientist* **36**: 536–44.

Weinberg, G., Godfrey, M., Rae, A. and Rhoads, J. 2008. A Real-Time Genetic Algorithm in Human-Robot Musical Improvisation. *Computer Music Modeling and Retrieval, Sense of Sounds*. Berlin: Springer, 351–9.

Whalley, I. 2009. Software Agents in Music and Sound Art Research/Creative Work: Current State and a Possible Direction. *Organised Sound* **14**: 156–67.

Whitelaw, M. 2004. *Metacreation. Art and Artificial Life*. Cambridge, MA: MIT Press.

Winkler, T. 1995. Strategies for Interaction: Computer Music, Performance, and Multimedia. *Proceedings of the 1995 Connecticut College Symposium on Arts and Technology*. http://www.brown.edu/Departments/Music/sites/winkler/research/papers/strategies_interaction_1995.pdf.

Wooldridge, M. 2009. *An Introduction to Multiagent Systems*. Chichester: Wiley.

Young, M. 2008. NN Music: Improvising with a 'Living' Computer. *Lecture Notes in Computer Science* 4969. Berlin and Heidelberg: Springer, 337–50.

https://doi.org/10.1017/S1355771811000094 Published online by Cambridge University Press